Homework 3: Exercise 1

Bernhard Sehorz

Problem: Write an MPI program in which two processes exchange a message of size N words. Repeat the experiment 10 times and calculate the average execution time. Use this program to measure communication bandwidth as a function of N on a single-local computer, and on two networked computers. Plot results. Obtain estimates for t_s and t_w for both cases. Discuss your results. To support your work, see the example "Sending in a ring (broadcast by ring)" on *http://www-unix.mcs.anl.gov/mpi/tutorial/mpiexmpl/.*) and the Lab exercise 4 in the file mpi_instructions.txt.

Solution: The code that was used to measure bandwith comes separately (`bandwidth.c`). It is a slightly modified version of the "Sending a ring"-example. It was run both on two networked computers and on a single host. The data obtained comes with the code (`resultsnet.txt` and `resultssingle.txt`). These data files where imported into Mathematica and graphically displayed (see `bsehorz_hw03.nb`).

The measurements consisted of sending messages of N integers (I took integers as words; an integer is usually 4 bytes), where N covered a range of 1 up to 100 000 with a step-size of ten. Each try for one message size was done ten times as requested in the assignment to average out network traffic peaks, CPU load peaks, etc. as can be seen from the code. All test were run in the computer room of the department of Scientific Computing.

In the case of the networked hosts, the maximum network bandwith is achieved pretty soon with message sizes of about 1 000 words. As I don't know anything about our network infrastructure, I cannot be absolutely sure about this result but assuming a 1-Gigabit-Ethernet connection, the maximum bandwidth of 90 million bytes per second seems very plausible. Of course the results obtained also depend on the current network traffic and CPU load. Since the tests were performed on thursday afternoon during the february semester break, I assume that the machines and network were not heavily in use. But one never knows...

In the case of a single computer, the maximum bandwidth is some orders of magnitude higher. (That means, it is (nearly exactly) 50 times as high: 4 500 compared to 90 million bytes per second.) It levels off at about 3 billion bytes per second around a message size of 70 000 words.

Considering t_w and t_s one can see from the plots of the execution time that for the single host $t_w = \frac{10^{-3}}{100000}$ (roughly 10 nanoseconds) and for the two networked hosts $t_w = 3\frac{10^{-2}}{100000}$ (roughly 300 nanoseconds). The startup cost t_s becomes almost irrelevant with message sizes increasing to 100 000 words. On a single processor it is about $2.5 \cdot 10^{-5}$ and two networked computers it is about $1.5 \cdots 10^{-4}$.